

TAB 1
TO EXHIBIT A
(PART 2)

Exhibit H

HIGH-DATA-RATE WIRELESS LAN
FOR TACTICAL MULTIMEDIA NETWORKS

Final Report

under contract

from U. S. Army Communications and Electronics Command

by

Jeffrey H. Fischer
Iwen Yao
Kendrick R. Bennett
John H. Cafarella

This document contains MICRILOR proprietary information. Distribution limited to U. S. Government agencies only. Requests for this document shall be referred to U. S. Army communications and electronics command.

MICRILOR, Inc.
17 Lakeside Office Park
Wakefield, MA 01880

CONTENTS

1	BACKGROUND	3
1.1	Motivation	3
1.2	High Data Rates in Multipath	4
1.3	Walsh-Function Signaling	5
1.4	Walsh Function Correlation Processors	7
2	IMPLEMENTATION APPROACH	7
2.1	Charge-Coupled Device Implementation	8
2.2	Digital Implementation	13
2.3	Receiver Concepts for the Digital Approach	15
2.4	Phase I Demonstration	16
3	CHIP DESIGN	16
3.1	Achieving High Speed	18
3.2	Clocking	21
3.3	Architecture for full I and Q processing	22
3.4	Chip Layout	22
3.5	Board Design	24
4	SIMULATION	25
5	REFERENCES	

APPENDICES

A1	MATLAB Code for Sub-Optimal I & Q Combining Analysis
A2	Schematics and Chip Layout
A3	Board Layout and Routing
A4	MATLAB Code for Simulation Program

1 BACKGROUND

1.1 Motivation

The requirements for high-data-rate connectivity in tactical Army deployments continues to exceed the rates available using emerging communications systems; these systems, which normally address the battlefield survivability problem by using considerable spectrum spreading, offer relatively low data rates and they also are very expensive. An alternate architecture for achieving the desired high-rate connectivity is to reduce the amount of protection obtained through spread-spectrum signaling, and to substitute a combination of antenna directivity and absorptive propagation behavior. On the other hand, it is not desirable to completely remove spread-spectrum signaling.

The use of absorption can provide tremendous enhancement of survivability by denying detection by the enemy's long-range surveillance assets; however, the use of absorption actually increases the detectability for an unintended receiver at nearer ranges. Thus, by including a modest amount of spreading, a hybrid approach to reduced probability of detection (RPD) can be effected which offers excellent protection against sophisticated stand-off surveillance and also against simpler intercept threats which might be closer to friendly locations. Another benefit of spread-spectrum signaling is the reduction of fading and/or intersymbol interference (ISI). Since each of these effects could potentially require increased transmitter power if spread spectrum were not used, there is a secondary enhancement of survivability besides the conventional RPD due to spreading. The problem, then, is to employ spread-spectrum signaling in order to gain its advantages without using up too much of the available bandwidth or greatly increasing the equipment cost.

1.2 High Data Rates in Multipath

The principle barrier to high-rate communications in a tactical ground-to-ground environment is the multipath nature of the channel. When transmitting binary waveforms whose duration is shorter than the spread in multipath delays, the channel response to each data bit overlaps that of other bits; this effect is called intersymbol interference (ISI). In a typical short-range (<1 km) link in open country (not downtown) this multipath spread might be about 1 μ sec, which would imply a maximum data rate of 1 Mbit/sec for binary signaling. (This multipath spread would be somewhat reduced in the network architecture based upon microwave links with modest antenna gain because the longest delays often correspond to angles far from bore-sight.) One effective means

MICRILOR PROPRIETARY: subject to restrictions on cover page.

for overcoming this limitation is to employ a higher-level signaling alphabet, i.e. send a "symbol" waveform whose duration is longer than the multipath spread, but which conveys multiple bits per symbol.¹ A natural set of orthogonal waveforms for "M-ary signaling" is the so-called Walsh functions;² these can be digitally generated and readily processed.

The use of M-ary orthogonal signaling is normally prohibited by the narrowband channel allocations available; in order to convey n bits per symbol the bandwidth required is $m=2^n$ times the symbol rate, and this exponential scaling rapidly expands the bandwidth. When it is intended, however, to employ spread-spectrum in the form of direct sequence, then the expansion due to M-ary signaling can be made to approach the spread-spectrum bandwidth without burden. Thus, we see that the combination of orthogonal signaling with DS spread-spectrum is a very natural approach to high-data-rate communication in multipath. In fact, a further attractive feature of this approach is that, when the multipath spread does occasionally exceed the symbol duration, then the use of changing codes for spreading results in randomization of the effects of ISI, which would enable coding to be employed to help overcome this residual ISI.

1.3 Walsh-Function Signaling

Figure 1 shows the probability of correctly demodulating a 1024-bit packet using DPSK spread-spectrum signaling and also using M-ary spread-spectrum with 2, 4 and 8 bits per symbol. It must be stressed here that the use of orthogonal signaling is a form of coding; if the DPSK modulation were combined with coding, it would also shift to a lower required S/J ratio.

Figure 2 gives the same curves on an expanded scale. Note that the use of 4-ary orthogonal (2 bits/symbol) is essentially equivalent to DPSK, and that a higher

signaling alphabet can be used to considerable effect. For an uncoded link, the above curve imply that the processing gain achieved by combining orthogonal signaling with DS spread spectrum is larger, for a fixed bandwidth. Thus, a transmission




Figure 1 - Probability of correctly demodulating a 1024-bit packet.

MICRILOR PROPRIETARY: subject to restrictions on cover page.

using 4 bits/symbol and a DS bandwidth of 80 MHz would exhibit the same nominal processing gain as a transmission using DPSK and a DS bandwidth of about 130 MHz!

It is important here to note the selection of Walsh functions here as the orthogonal set. In JTIDS a rotation of the DS code known as cyclic code-shift keying (CCSK) was used; this produced, at the processor output, the equivalent of pulse-position modulation (PPM). In a multipath environment the PPM-like behavior of CCSK can be very detrimental because the delay discriminant used for demodulation can be directly confused with multipath.

Thus, it is preferable to use some other orthogonal set. If the bandwidth expansion due to the orthogonal signaling is much smaller than the DS spreading, then one might use essentially any orthogonal set; in this case a convenient choice would be a set of tones (i.e., M-FSK) which could be demodulated using FFT techniques after the spreading code is stripped off. However, when the spreading due to the orthogonal signaling approaches that of the DS, then the actual bandwidth will depend upon both the DS and orthogonal modulation. It is interesting that CCSK and the use of Walsh functions both have the property that the bandwidth required for orthogonal signaling can equal the DS bandwidth without further expansion of the signal bandwidth over that of the DS. Thus, we see that for confining the bandwidth and performance in multipath, the Walsh functions are a natural choice.




Figure 2 - Expanded view of demodulation performance.

1.4 Walsh Function Correlation Processors

In this section, we assume that the reference DS code is properly aligned with the received signal, and that the two are multiplied, and then low-pass filtered to support processing the highest Walsh function. (This generally may be lower than the chip rate.) This has the effect of "stripping off the DS code, leaving only the Walsh function modulation. The discussion below considers baseband processing; in an actual implementation the signal would be converted to in-phase and quadrature channels, and the baseband processing described below would be performed in

MICRILOR PROPRIETARY: subject to restrictions on cover page.

each channel, followed by envelope combining of corresponding amplitudes before data decision.

One approach to forming the required correlations would be to use a fast Walsh transform; this is awkward, and would become attractive only if the order of the Walsh functions became large. Another approach to realizing the correlator is to compute the coefficients in a pipelined architecture, as shown in Fig. 3, utilizing the structure of the Walsh functions to reduce the clock rate of successive stages which are increasing in the number of parallel cells, thus keeping the power per stage nearly constant. A basic cell accepts input samples at rate f_{in} and generates two output samples at rate $f_{in}/2$: one output sample is the sum of the current pair of input samples and the other is the difference of these two input samples. The basic cell and a tree for decoding the first 8 Walsh functions is shown below. Notice that a particular path through this tree corresponds to multiplying the input sequence by a specific combination of square-wave functions and summing the result. There is a natural conservation of power in that each successive stage has twice the hardware but runs at half the clock rate.




Figure 3 - Architecture for Walsh coefficient computation.

2 IMPLEMENTATION APPROACH

We chose to implement the processor in a field programmable gate array for expediency. Also, any follow-on effort would benefit from this as re-usable circuits. Higher levels of integration are available for the higher density devices that a follow-on effort would require. We also examined the possibility of a more efficient implementation for a production version using CCDs. Although the CCD technology we found available would probably support both the synchronization and demodulation functions, a custom design would have to be implemented. Such a design is beyond the scope of the phase I work and should be re-examined in a follow-on effort. The CCD would take discrete-time analog samples at its input, perform all the signal processing and produce digital data decisions at its output. The following section summarizes the issues uncovered in researching the CCD approach.

2.1 Charge-Coupled Device Implementation

Since the CCD input is a capacitive storage cell, the analog signal is inherently sampled as the CCD input is clocked. This operation is similar to the sampling capacitor in a sample and hold circuit. Alternatively, a digital approach would require high speed D/A converters that are very costly. The input linearity would support far more processing gain than necessary for this project and at a sampling rate of 50 Msps. This rate includes a factor of two oversampling to prevent excessive straddling loss. Higher rates may be possible by operating two in parallel and staggering the clocks. The signal is processed with analog precision against a binary digital reference.

These devices are single-phase clocked with a separate loading clock for the reference. The reference is double buffered so they would support continuously changing codes between symbols.

Since the processor can handle a chipping rate of 25 MHz, a Walsh demodulator could be configured out of CCD cells. If the direct sequence code is of higher bandwidth, that could be stripped off in a pre-correlator once chip timing has been established during the synchronization cycle. Establishing synchronization could also be performed in the CCD, either by a reconfiguration of the demodulation CCD channels or by a separate CCD channel operating in parallel. If the same chip rate is preferred for both synchronization and demodulation, and if it is higher than 25 MHz, parallel CCDs would have to be used combined with some multiplexing structure to implement a matched filter. Alternatively, recognizing that there is no baud rate requirement for synchronization, a longer symbol can be used with a lower

MICRILOR PROPRIETARY: subject to restrictions on cover page.

direct sequence code rate preserving the processing gain. The number of code chips in the matched filter can be further increased so as to provide a robust synchronization preamble. This way, for example, a spread spectrum code rate of 25 MHz could be used for synchronization, while a Walsh chip rate of 25 MHz could be used for data with a 75 MHz spread spectrum code under it that would be processed away in the pre-correlator.

A enticing advantage to CCDs is that all the support circuit can be integrated onto the same device as the storage cells. Since modern CCDs use dynamic-RAM-like storage, these devices have benefitted directly from achievements in the packing density and speed of DRAM technology.

Such circuits as the magnitude detections and 8-way comparisons discussed later are conceivably integratable directly on the device output which would greatly simplify the output processing. The cost of these devices is projected to be an order of magnitude lower than for digital matched filter chips.

2.2 Digital Implementation

2.2.1 Circuitry reductions for digital FPGAs

The FPGA technology has the advantage of very fast turnaround to obtain a working design. The tradeoff is that the speed and packing density have some limitations. The device used here is provided by XILINX, inc. and allows for the equivalent of several thousand digital gates. There are other practical limitations such as "bugs" in the FPGA compiler software that XILINX provides. Any follow on effort would probably center around an implementation in FPGAs as the time to achieve a working design is significantly shorter than either a semicustom chip, or a board of LSI logic. As a consequence, some effort was expended to analyze the losses incurred for several reasonable approximations that would reduce the gate count.

2.2.2 Study reduced number of input bits

The number of logic cells used in the device is proportional to the number of input bits that are taken for each sample. For the Walsh correlator section, the number of cells is $C \cdot (7 \cdot B + 13)$ per channel (I or Q) for an 8-ary alphabet. Where C is the number of cells used to process a single bit per stage, and B is the number of input bits. The factor of 7 and the residual 13 are functions of the number of stages. This results in about 165 cells for a 6-bit input. Reducing the number of input bits has an impact of

MICRILOR PROPRIETARY: subject to restrictions on cover page.

21 cells per bit! Fig. 4 is a plot of implementation loss versus the number of bits kept for a correlation process in a Gaussian channel³. It can be seen that a few bits helps a lot, but only if the signal is properly normalized. The more bits kept, the less stringent are the normalization requirements as seen by the broadening of the curves. We also considered the performance in a non-Gaussian channel and determined that the risk of substantial performance loss was not worth reducing the number of bits below 6 at the input (5 plus sign) for the Phase I effort.




Figure 4 - Implementation loss vs. noise normalization for various bits of quantization.

A tremendous payoff would occur if binary decisions could be made at the input instead of multibit quantization in the A/D converters. Besides a factor of three reduction in logic cells, the A/D converters would be eliminated! An ad hoc approach to determining the loss was used, recognizing the number of incorrect decisions at the input that would be required to cause an error at the output. To explain this approach we inspect the set of the first 8 Walsh functions as shown in Fig. 5. These binary waveforms would phase modulate a carrier, be transmitted over the air and, and be corrupted by noise. At the receiver the carrier would be removed and the results would be filtered to obtain analog, noisy, samples of this waveform at the transitions. Normally, these transitions would be correlated against by analog multiplication with the reference waveforms followed by analog summing. If, prior to correlation, the waveforms were first quantized into binary decisions, the correlation process reduces to adding up the number of agreements to each of the reference waveforms and comparing the results. If there are no decision errors at the input, then the correct reference correlator would have a value of 8 (for the 8-ary case) and the others would have 0 as they are orthogonal.

To see how an output error occurs, we can assume that two decision errors were made at the input. Instead of adding in the correct channel, these two decisions subtract from the total resulting in a value of 4 at the output. If there is an

MICRILOR PROPRIETARY: subject to restrictions on cover page.

incorrect channel that will have these two samples adding up instead of canceling, then it also will have a value of 4. When this occurs, the probability of error is 50%. We will say that an error occurs. Inspection of the waveforms will indicate that given any two Walsh chips in one waveform, there is at least one other waveform that has the reverse relative polarity. So that if the first waveform has chips 1 and 7 as ++, there is at least one other waveform with chips 1 and 7 as either +- or -+, and vice versa. This means that for any pair of errors, there is at least one waveform that will convert the pair of decisions that were adding up into one that cancels, or one that cancels into one that adds up. The result is that a value of 4 occurs in that channel, and thus an error. Computing the SNR that would be required to support binary decisions at the input with less than 2 errors results in about a 2.1 dB loss over the analog case.

To compute the loss, the Binomial probability of getting 2 errors of 8 trials is set equal to an operating point for 8-ary orthogonal. For example, for an output $P_e = 10^{-5}$, the input P_e is $6E-4$. For PSK demodulation, this input P_e requires a E_b/N_0 of the Walsh chip = 6.6 dB. For the analog case, the symbol E/N_0 for an output $P_e = 10^{-5} = 13.5$ dB which translates to a 4.5 dB E_b/N_0 for the Walsh chip.

The 16-ary case exhibits about a 3.25 dB loss. The loss based on comparisons of binary cases would be expected to decrease with an increased number of input decisions. Perhaps the m-ary signalling improves faster than the binomial decision process. Additionally, binary quantization is known to exhibit very poor performance in non-Gaussian interference. Since the Walsh bandwidth is of the same scale as the spread spectrum bandwidth,




Figure 5 - Walsh functions in Walsh order.

MICRILOR PROPRIETARY: subject to restrictions on cover page.

the direct sequence can not be used to make the interference appear nearly Gaussian. Therefore, binary quantization was abandoned for now. It is recommended that additional analytical effort be expended to make a final judgement on this issue.

2.2.3 Suboptimal I and Q combining

The number of logic cells required to implement the I and Q channel combining is substantial. The amplitudes are ideally squared and summed. The final square root is not necessary unless dynamic range is a problem. Since these operations are hardware intensive, unless performed in a DSP microprocessor (at slower rates), approximations are typically used. Determining the largest of the I or Q channels and adding to it 1/4 of the smallest is a fair approximation. This requires a byte-wide comparator, multiplexer and summer per correlator channel I and Q pair. An improved approximation is to take the largest plus 3/8 of the smallest. This requires an additional adder. The byte-wide logic would occupy about as many logic cells as the 8-way Walsh correlator itself. Thus we analyzed some other notions of approximations. These results are shown in the curves of Fig. 6. The MATLAB code to do the analysis is included in appendix A1. The curves compare the ideal performance with simple algorithm of summing the I and Q channel linear magnitudes, and with selecting the largest of the I or Q channels to use in the comparison. The 3 dB loss seen in the curves was deemed too much to incur.

MICRILOR PROPRIETARY: subject to restrictions on cover page.




Figure 6 - Performance of several approaches to sub-optimal combining of I and Q.

MICRILOR PROPRIETARY: subject to restrictions on cover page.

2.3 Receiver Concepts for the Digital Approach

Fig. 7-a and 7-b show two receiver concepts for the Walsh demodulator. The Bandpass precorrelation of Fig. 7-a would be performed by a mixer and SAW bandpass filter which would provide the square-shaped impulse response that would filter the input waveform to the Walsh bandwidth. This bandwidth would then be sampled in I and Q channels by the slow(er) A/D converters. The digital conversions would be passed through the Walsh demodulator that would perform the 8-way correlation, the resulting I and Q channels would be combined and the largest output would indicate the most likely data symbol transmitted. The difference between this diagram and that in Fig. 7-b is that the latter does not attempt to pre-correlate the analog signal. The spread spectrum is stripped off and the band-pass filter is broad. The A/D converters must be faster to operate on the full spread spectrum bandwidth which in general can be wider than the Walsh bandwidth. The correlator pre-summing prior to Walsh demodulation must be performed digitally at the full spread-spectrum code rate. This approach saves the investment in a SAW filter and might make sense in prototype quantities.




Figure 7 - Correlation approaches; a) bandpass pre-correlation, b) baseband pre-correlation, and c) the approach used in the demonstration hardware.

MICRILOR PROPRIETARY: subject to restrictions on cover page.

2.4 Phase I Demonstration

It was assumed that the purpose of implementing hardware in a phase I study is to demonstrate the feasibility of higher risk items, thereby removing this risk from a follow-on effort. The detailed design of the parts that have not been previously tried were deemed the high risk items. Specifically these items were the 8-way Walsh correlator and the 8-way compare. While the I and Q combining is substantial circuitry, it had been done several times in the past by MICRILOR and by others, and is considered well understood in terms of the errors involved, the amount of circuitry involved, the timing, the SNR losses, and the dynamic range requirements.

It was also assumed that the distinctions between the known-phase and the unknown-phase demodulation was well understood. Since I and Q channels would be required for demodulation of a signal of unknown phase, there would be a substantial savings if we assumed a known phase (as if the phase were acquired through a tracking loop, for example). This assumption would save one channel and the I and Q combining circuits. Thus we would demonstrate a single channel 8-way Walsh correlator, and an 8-way magnitude compare (after stripping the polarity of the output signal). Furthermore, it was decided that the RF circuits and spread spectrum circuits have been demonstrated at MICRILOR and other places many times and constructing such circuits would not contribute to reducing the risk of a follow-on effort. All the information that these circuits could generate to exercise the circuits that were being tried could be generated by computer and fed into the Walsh demodulator from memory.

Therefore, for the demonstration, a circuit board was constructed with high speed RAM that could be downloaded from the computer that would load the data into the demodulator at the full design clock rate. The 10 Mbps data is read out of the demodulator into a logic analyzer for examination. The block diagram of this design is shown in Fig. 7-c. The Walsh correlator, sign stripping, 8-way compare, and 8 to 3 decoder are all included on a single XILINX FPGA along with clocking and control circuits. The circuit board will also have two D/A converters for testing. A version of the XILINX design was made with just a pair of 4-byte to 1-byte output muxes so the user could examine any pair of channel values on an oscilloscope.

MICRILOR PROPRIETARY: subject to restrictions on cover page.

3 CHIP DESIGN

3.1 Achieving High Speed

The chip architecture was developed from the basic building block and architecture in fig 3. This architecture is efficiently achieved in a XILINX FPGA as it is register intensive and pipelined. An advantage to this particular architecture is that while each stage doubles the circuitry from the previous stage, the later stage is only required to operate at half the speed. This balances out the power distribution on the chip, and, as well, means that the most speed-intensive circuits do not require as much routing. Thus the higher speed input stages can use optimized routing and the last stage, which has the highest adder fanout and thus requires much more routing channels, can be poorly placed and not optimized since it runs relatively slowly.

The first stage adder must operate on a six-bit twos-complement input with sign extension to produce a seven bit output. It must operate at 26.67 MHz clock rate to support 10 Mbps 8-ary signalling. In spite of using the fastest chip available in FPGA technology, these constraints required that this part of the design be placed and routed on the chip as a separate block. A special adder circuit was designed that used some amount of carry lookahead so that the carry levels did not create excessive delay. Once this section was complete and was simulated, it was locked in and the next stage in the pipeline was created. This was placed and routed as a add-on to the high-speed input stage. Similarly the third stage was added. The 8-way (1 channel only) Walsh correlator used about half the chip or 168 logic cells. The sign stripping and 8-way compare compiled separately into another 76 cells or about half as much as the Walsh correlator. Due to a "bug" in the XILINX software we experienced difficulties in merging these two designs. (This technology is only a few years old and is still maturing). As a result, the final design is borderline on speed and we anticipate faster FPGA chips will be available in the next few months to cure this.

A block diagram of the final chip design is shown in Fig. 8. As can be seen, the 8-way compare is performed in a tree-like structure that takes the inverse form of the Walsh correlator fan-out expansion.

MICRILOR PROPRIETARY: subject to restrictions on cover page.

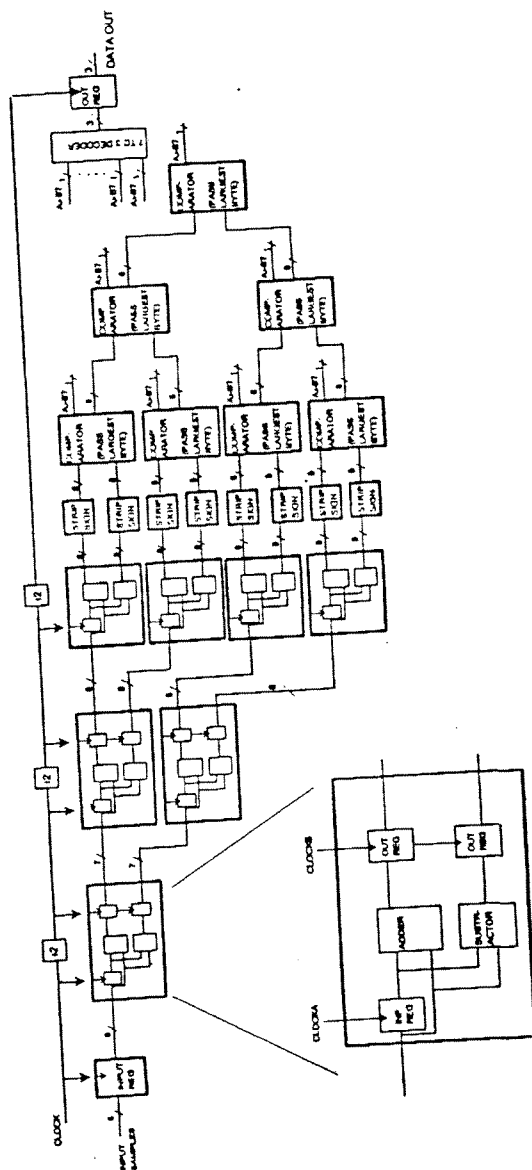


Figure 8 - Final chip design.

MICRILOR PROPRIETARY: subject to restrictions on cover page.

3.2 Clocking


3.2.1 Pipeline Delay

With the clock scheme shown, the pipeline delay is 15 input clock cycles for an 8-ary symbol. Thus there is nearly a full symbol of delay, after the symbol of concern is completely sampled, before the corresponding data appears at the output. (The minimum delay for a parallel correlator structure would be 8 input clocks).

3.2.2 Maximum Clock Rate

After a pair of input samples are loaded to the first Add/Sub block, the result must be latched out by the very next sample. Thus, even though the input stage operates on a pair of samples at a time, the result must be ready in a single sample time. Similarly, the second stage operates on a pair of output samples of the input Add/Sub block. This spans four input sample times. However, the result must be latched out to the next stage before its next input sample arrives and thus it must operate at one half the input rate. The third Add/Sub block has been modified so that the output register is placed after the 8-way compare. This eliminates an extra frame delay. This third block and the entire output circuit must operate at one fourth the input rate. If double buffering were used at the input to each stage, these operating rates would be halved as an extra set of samples would be stored while awaiting the computation. The number of registers would overwhelm the design and operating at twice the rate is a preferred approach.

There is typically some skew that is expected between the arrival of the clocks to the latches at the inputs and outputs of each stage. The difference between the latest arrival time of the input clock to the earliest arrival of the output clock, if positive, must be added to the maximum delay path of the logic to determine clocking rate. Excess delay may be added to the output clock to increase the speed, provided care is taken that minimum hold times are met for the shortest logic delay path that could occur.



MICRILOR PROPRIETARY: subject to restrictions on cover page.




Figure 10 - Preferred partitioning of full design.

MICRILOR PROPRIETARY: subject to restrictions on cover page.

3.3 Architecture for full I and Q processing

The pipelined architecture described in Fig. 8 must be duplicated for the Q channel, and the I and Q combining must be added along with the 8-way compare. A block diagram for the entire design is shown in Fig. 9. The quadrature-channel combining is assumed to be a $X + 1/4 Y$ technique, and therefore the sign removal is still necessary on each channel. Also note that if either the I or Q channel have a maximum value of 8 bits, the magnitude also has an 8 bit maximum as it will lie on a circle of that radius.

The most straight-forward way of fitting this design into FPGA chipss to put a single Walsh correlator in each chip, bringing the inputs to the I and Q combining circuits onto a third chip that also contains the 8-way compare. This approach requires coming off each correlator chip with 64 lines (8 bits by 8 channels) and into the combiner chip with 128 lines; This amount of I/O is stressful to the design cost and reliability. A better approach would be to slice the design up as in Fig. 10. Inspection of the block diagram shows that the upper and lower correlators in each channel are not connected until the final comparison. It is the upper correlators of each channel that must be connected in the middle of the data-stream. It isn't until each set of I and Q correlators are paired down in the comparison tree that they must join. In Fig. 10 the circuits above the dashed line are in one device and those below the dashed line are in the another. Only 24 data lines connect between the two chips.

In the 8-ary design as shown in the Fig. 10, after the first stage, 7 lines in each of the I and Q channels are brought to the other chip, and 4 complex correlators are processed per chip. The 4 I and 4 Q correlators are identical to those in Fig. 9, as if they were separated into I and Q chips, and after the 4 I and Q results are combined, they are paired down to a single largest value in a 4-way compare. This 8-bit value (along with 2 decoding lines) is exported to the other chip where it is compared to the result from the other half of the compare tree and the largest of the 2 values is decoded. The result is a three bit word producing the data for the demodulated symbol, just as in the single channel, phase I demonstration.

The next higher family of devices may be required for this level of integration (conversion of the design is automatic). Or, the design may be further subdivided by parsing each section similarly and four smaller chips may be used. Similar parsing may be used for higher order alphabets. It may be necessary to separate out a front end source that then feeds individual correlator banks. The tree structure still allows reduction of

MICRILOR PROPRIETARY: subject to restrictions on cover page.

the fanout on the same chip. For improved operating speed, necessary in higher alphabets, the number of signal lines between chips can be traded against an optimized high speed front end.

This approach allows use of the multiple chips of FPGA technology in a follow-on effort in order to expedite the construction of a full receiver. In hindsight, this parsing would have simplified the phase I effort as well as it would have avoided the need to work around the compiler bug.

3.4 Chip Layout

After many place and route iterations, the chip design was finalized and exercised in a logic simulator for design correctness and timing. The resulting schematics and chip layout are included in Appendix A2. The chip layout shows the routing channels occupied between the logic cells. The bent lines occur when a line crosses routing channels. Also included is the test layout that for the design that allows looking at the output channels on the oscilloscope.

3.5 Board Design

The circuit board schematic is shown in Fig. 11. It comprises a high speed static RAM, a XILINX chip, two high speed counters for addressing the RAM, a download connector for the PC to load the RAM and a tri-state buffer for this port, two D/A converters, and loading circuitry for the XILINX. (The XILINX routing and configuration information is actually dynamically stored and gets automatically loaded each time the power is turned on). For reference, the Board layout and routing is shown in Appendix A3. The samples that are downloaded to the RAM are generated in the simulation program where noise and interference is injected.

The D/A converters allow examination of the magnitude of the 8 correlator outputs over time on an oscilloscope to ensure correct operation. When operating with the 8-way compare, the outputs are generated on 3 signal lines representing the 3 bits that change each symbol. This can be clocked into a logic analyzer which can be triggered off the terminal count of the counters that drive the RAM address bus.

MICRILOR PROPRIETARY: subject to restrictions on cover page.




Fig 12 DIGITAL BOARD BLOCK DIAGRAM

Figure 11 - Circuit board schematic.

MICRILOR PROPRIETARY: subject to restrictions on cover page.

4 SIMULATION

A program was written with MATLAB software to generate the input samples to the XILINX chip. The synthesized input samples were created by simulating the transmitter, the communication channel, and the receiver front end and pre-correlator. After testing the initial code, it was recognized that this program could readily be modified to simulate the Walsh demodulator. Thus, the performance of a Walsh demodulator under various channel conditions could be studied. Since this simulation was a fortunate by-product of the original program, and since it was deemed an important tool for this work, considerable effort was then expended to improve the run-time of the program so that Monte Carlo simulations could be performed. The effort resulted in a 50-times speed improvement, making the simulator capable of demodulating thousands of symbols in under a minute on a 486-based PC. Thus, useful Monte Carlo runs can now be made. Multipath profiles can be created, noise and interference and Doppler can be injected, and non-idealities such as finite input-sample word size can be simulated.

Closed-form analysis of a demodulator is typically either performed for a benign additive white Gaussian noise channel or several selected but often contrived cases of corrupted channel. Usually, the compromises made to the problem in order to complete the analysis reduces the impact of the results. With this simulation tool, nearly any channel can be created and the performance observed. While this is not a substitute for the experience gained by actual usage, it can provide an intuitive feel for how the communications link will perform.

The following functions are performed in by the simulation program:

- Generation of random data;
- Walsh modulation of the data;
- Multiplication by spread spectrum code and carrier;
- Addition of multipath;
- Addition of channel noise and interference;
- Stripping off the carrier;
- Multiplication by reference spread spectrum code;
- Pre-correlation over 3 chips;
- Truncation of result to demodulator input word size;
- Save data for digital board to exercise XILINX chip;
- Perform 8-way correlation in I and Q channels;
- Results for I channel alone and for combined I and Q;
- Perform 8-way compare on each set of results;

MICRILOR PROPRIETARY: subject to restrictions on cover page.

The simulation program is supplied in appendix A4.

Simulation runs for a single path were made to check the simulator against known analytical results. They checked to within 0.5 dB. Then, a worst case multipath scenario was simulated whereby two equal-power multipath signals were swept in delay in 5 ns increments over a range of several symbols. At each delay, the P_e was measured. The SNR was then increased and the run performed again until the P_e of the original single path run was matched across most of the delay values. It was found that for a single path operating point of $P_e=10^{-3}$, many of the delay values were unaffected, and a 3 dB increase in SNR was sufficient to remove the affect of the equal power multipath for most of the delays. The few remaining paths required a couple more dB in SNR. Had this worst-case scenario been used in a narrowband system, the error rate would be unrecoverable at many delays as there would be 100% cancellation. This equal-power two-path scenario is much worse than the often quoted Rayleigh fading probability density.

Runs were also made for a two-path operating point of $P_e=10^{-5}$. An increase of 5 dB in SNR was required to recover the P_e for many of the delay values and 8 dB was sufficient for the rest. It is often the case in that operating points with lower P_e are more sensitive and will suffer greater loss.

A copy of the Matlab code for the simulation has been made available to MIT Lincoln Laboratory because of its interest in developing CCD technology for Walsh-orthogonal spread-spectrum signaling. This collaboration should result in more extensive running of simulations of interest to the Army than has been possible to date; MICRILOR has only a modest-speed workstation (i.e. an IBM 486) and limited resources under Phase I funding.

5 REFERENCES

1. J. M. Wozencraft and I. M. Jacobs, Principles of Communication Engineering. New York: Wiley & Sons, 1965, pp. 257-266.
2. K. G. Beauchamp, Applications of Walsh and Related Functions. New York: Academic Press, 1984.
3. J. H. Cafarella, "Advanced SAW-Based Signal Processing for Packet Communication," (Invited), MILCOM '82 Proc., vol. S, Progress in Spread-Spectrum Communication, New York: IEEE, 1982, pp. 10.6-1 - 10.6.7.

MICRILOR PROPRIETARY: subject to restrictions on cover page.